

# The Color of the Cat is Gray: 1 Million Full-Sentences Visual Question Answering (FSVQA)

Andrew Shin, Yoshitaka Ushiku, Tatsuya Harada

The University of Tokyo  
7 Chome-3-1 Hongo, Bunkyo  
Tokyo 113-8654, Japan

## Abstract

Visual Question Answering (VQA) task has showcased a new stage of interaction between language and vision, two of the most pivotal components of artificial intelligence. However, it has mostly focused on generating short and repetitive answers, mostly single words, which fall short of rich linguistic capabilities of humans. We introduce Full-Sentence Visual Question Answering (FSVQA) dataset ([www.mi.t.u-tokyo.ac.jp/static/projects/fsvqa](http://www.mi.t.u-tokyo.ac.jp/static/projects/fsvqa)), consisting of nearly 1 million pairs of questions and full-sentence answers for images, built by applying a number of rule-based natural language processing techniques to original VQA dataset and captions in the MS COCO dataset. This poses many additional complexities to conventional VQA task, and we provide a baseline for approaching and evaluating the task, on top of which we invite the research community to build further improvements.

## Introduction

The research community in artificial intelligence (AI) has witnessed a series of dramatic advances in the AI tasks concerning language and vision in recent years, thanks to the successful applications of deep learning techniques, particularly convolutional neural networks (CNN) and recurrent neural networks (RNN). AI has moved on from naming the entities in the image (Mei et al. 2008; Wang et al. 2009), to describing the image with a natural sentence (Vinyals et al. 2015; Xu et al. 2015; Karpathy and Li 2015) and then to answering specific questions about the image with the advent of visual question answering (VQA) task (Antol et al. 2015).

However, current VQA task is focused on generating a short answer, mostly single words, which does not fully take advantage of the wide range of expressibility inherent in human natural language. Just as we moved from merely naming entities in the image to description of the images with natural sentence, it naturally follows that VQA will also move towards full-sentence answers. One way to tackle this issue would be to apply appropriate linguistic rules *after* a single-word answer is generated. However, previous works in natural language processing field have demonstrated that data-driven response generation achieves better performance than rule-based generation (Ritter, Cherry, and Dolan 2011). In other words, it is more efficient to provide data only once

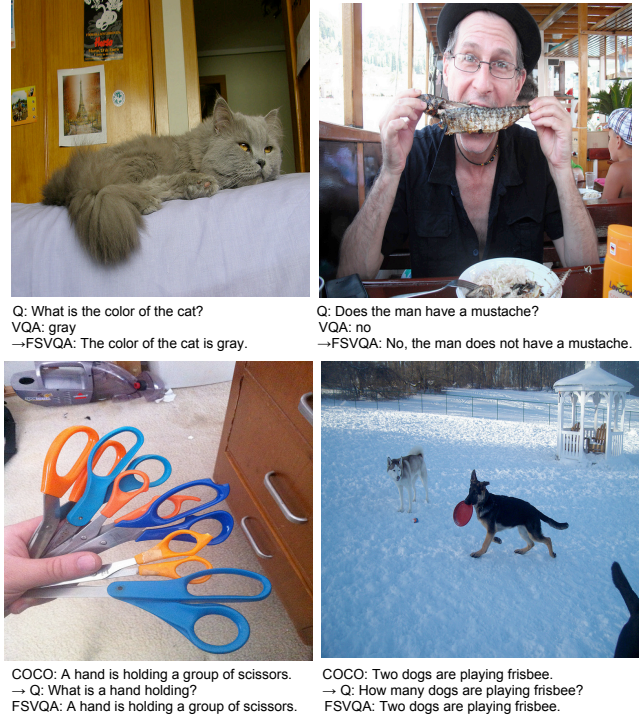


Figure 1: Examples of FSVQA dataset. Top row shows examples of full-sentence answers generated by converting VQA questions and answers. Bottom row shows examples of questions and answers generated by converting MS COCO captions.

for pre-training than to parse and tag the text every time to apply universal rules. In addition, training with full-sentence answers provides an opportunity for the learning of complex morphological transformations along with visual and semantic understanding, which cannot be done with manual application of rules.

Learning and generating full-sentence answers will inevitably increase the number of distinct answers at an exponential scale; for example, thousands of samples with simple identical answer “yes” will be further divided into “yes, the color of the car is red,” “yes, the boy is holding a bat,” etc. Indeed, our FSVQA dataset contains almost 40 times more

answers that are unique than the original VQA dataset. This poses additional challenges on top of original VQA task, since now it not only has to come up with the correct answer, but also has to form a full sentence considering how the words are conjugated, inflected, and ordered.

We introduce Full-Sentence Visual Question Answering (FSVQA) dataset, built by applying linguistic rules to original VQA dataset at zero financial cost. We also provide an augmented version of FSVQA by converting image captions to question and answers. We examine baseline approaches, and utilize complementary metrics for evaluation, providing a guideline upon which we invite the research community to build further improvements.

Our primary contributions can be summarized as following: 1) introducing a novel task of full-sentence visual question answering, 2) building a large, publicly available dataset consisting of up to 1 million full-sentence Q&A pairs, and 3) examining baseline approaches along with a novel combination of evaluation metrics.

## Related Work

A number of datasets on visual question answering have been introduced in recent years (Malinowski and Fritz 2014; Ren, Kiros, and Zemel 2015), among which (Antol et al. 2015) in particular has gained the most attention and helped popularize the task. However, these datasets mostly consist of a small set of answers covering most of the questions, and most of the answers being single word. Our FSVQA dataset, derived from (Antol et al. 2015), minimizes such limitation by converting the answers to full-sentences, thus widely expanding the set of answers.

(Fukui et al. 2016) proposed multimodal compact bilinear pooling (MCB) to combine multimodal features of visual and text representations. This approach won the 1st place in 2016 VQA Challenge in real images category. (Saito et al. 2016) proposed DualNet, in which both addition and multiplication of the input features are performed, in order to fully take advantage of the discriminative features in the data. This method won the 1st place in 2016 VQA Challenge in abstract scenes category.

(Yang et al. 2016) was one of the first to propose attention model for VQA. They proposed stacked attention networks (SANs) that utilize question representations to search for most relevant regions in the image. (Noh and Han 2016) also built an attention-based model, which optimizes the network by minimizing the joint loss from all answering units. They further-proposed an early stopping strategy, in which overfitting units are disregarded in training.

(Lu et al. 2016) argued that not only visual attention is important, but also question attention is important. Co-attention model was thus proposed to jointly decide where to attend visually and linguistically. (Kim et al. 2016) introduced multimodal residual network (MRN), which uses element-wise multiplication for joint residual learning of attention models.

Most of the works above limited the number of possible answers, which was possible due to a small number of answers covering the majority of the dataset. Our

FSVQA dataset imposes additional complexity to existing approaches by having a much larger set of possible answers, in which no small set of labels can cover the majority of the dataset.

## Dataset

Collecting full-sentence annotations from crowd-sourcing tools can be highly costly. We circumvent this financial cost by converting the answers in the original VQA dataset (Antol et al. 2015) to full-sentence answers by applying a number of linguistic rules using natural language processing techniques. Furthermore, we also provide an augmented version of dataset by converting the human-written captions provided in the MS COCO (Lin et al. 2014). We generated questions with a set of rules, for which the caption itself becomes the answer. Both versions of FSVQA dataset along with the features used in our experiment, as will be described in the Experiment Section, are publicly available for download. Note that, for both versions, only train and validation splits are provided, since test splits are not publicly available. Also, we only provide open-ended version, and do not provide multiple choice version.

## Converting VQA

VQA dataset comes with 10 annotations per question, and we chose one annotation per question that has the highest frequency as the single answer for corresponding question. If par, one annotation was randomly selected.

VQA dataset mainly consists of three categories of questions; yes/no, number, and others. Table 1 summarizes the general conversion rule for generating full-sentence answers for each category, along with examples. Part-of-speech tag notation follows that of PennTree I Tags (Marcus et al. 1994), except NP and VP refer to parse tree instead of than part-of-speech.  $tense: \mathbb{V} \rightarrow \mathbb{T}$  returns the tense of the input verb,  $conjug: \mathbb{V} \times \mathbb{T} \rightarrow \mathbb{V}$  conjugates the input verb to the input tense, where  $\mathbb{V}$  is a space of verbs of all forms, and  $\mathbb{T}$  is a set of tenses such that  $\mathbb{T} = \{past, present, future, past\ perfect, present\ perfect, future\ perfect\}$ , except it returns the input as is if the input is of JJ tag.  $negate: \mathbb{V} \rightarrow \mathbb{V}$  negates the input verb of given tense, and  $replace(A,B)$  substitutes A by B. Parentheses indicate an optional addition, while A/B indicates an insertion of one of the two sides depending on the question. To briefly illustrate the process, these general conversion rules substitute the question phrase with the answer, and reorder the sentence with appropriate conjugation and negation.

While these general rules cover the majority of the questions, some types of questions require additional processing. For example, conditional statements with “if,” or selective statements such as “who is in the picture, A or B?,” can be handled by disregarding the sub-clauses and applying the rules to the main clause. Also, since VQA dataset consists of human-written natural language, it inevitably contains variations encompassing colloquialism, typos, grammar violations, and abridgements, which make it difficult to apply any type of general conversion rule. We either manually modify them, or leave them as they are.

Table 1: General conversion rules for generating full-sentence answers.

Type	Rule (Q→A)	Question	Ans.	Converted Ans.
yes/no	VB1+NP+VB2/JJ? →“Yes,”+NP+conjug(VB2/JJ,tense(VB1)) <b>or</b> , “No,”+NP+negate(conjug(VB2/JJ,tense(VB1)))	– <i>Did he get hurt?</i> <i>Is she happy?</i>	– yes no	– <i>Yes, he got hurt.</i> <i>No, she is not happy.</i>
	MD+ NP+VB? →“Yes,”+NP+MD+VB <b>or</b> , “No,”+NP+negate(MD)+VB	– <i>Will the boy fall asleep?</i> <i>May he cross the road?</i>	– yes no	– <i>Yes, the boy will fall asleep.</i> <i>No, he may not cross the road.</i>
number	“How many”+NP+is/are+EX?→EX+is/are+ans+NP	<i>How many pens are there?</i>	2	<i>There are 2 pens.</i>
	“How many”+NP1(+MD)+VB(+NP2)? →ans(+MD)+VB(+NP2)	– <i>How many people are walking?</i>	– 3	– <i>3 people are walking.</i>
	“How many”+NP1+VB1/MD+NP2+VB2? →NP2+(MD+VB2)/conjug(VB2,tense(VB1))+ans+NP1	– <i>How many pens does he have?</i>	– 4	– <i>He has 4 pens.</i>
	WP/WRB/WDT+“is/are”+NP? → NP+“is/are”+ans. WP+NP+VP? → ans.+VP WDT+NP+VP(+NP2)?→ans.(+NP)+VP(+NP2) WP/WDT+MD+VB?→ans.+MD+VB WP/WDT+MD+NP+VB?→NP+MD+VB+ans. WP/WDT+VP(+NP)?→ans.+VP(+NP) WP/WDT+VB1+NP+VB2? →NP+conjug(VB2,tense(VB1))+ans.	<i>Who are they?</i> <i>What food is on the table?</i> <i>Which hand is holding it?</i> <i>Who would like this?</i> <i>What would the man eat?</i> <i>Who threw the ball?</i> – <i>What is the man eating?</i>	students apple left dog apple pitcher – apple	<i>They are students.</i> <i>Apple is on the table.</i> <i>Left hand is holding it.</i> <i>Dog would like this.</i> <i>The man would eat apple.</i> <i>Pitcher threw the ball.</i> – <i>The man is eating apple.</i>

Table 2: General conversion rules for converting captions to questions.

Type	Rule (C→Q)	Caption	Question
yes/no	NP→“Does it look like”+NP	<i>A man.</i>	<i>Does it look like a man?</i>
	NP1 + VB (+NP2) →conjug(“do/be”,tense(VB))+NP1+conjug(VB,present)(+NP2)?	– <i>A dog jumped.</i>	– <i>Did a dog jump?</i>
	NP1+MD+VB(+NP2)→MD+NP1+VB(+NP2)? EX+“is/are”+NP→“is/are”+EX+NP?	<i>A boy would hit the ball.</i> <i>There are cats.</i>	<i>Would a boy hit the ball?</i> <i>Are there cats?</i>
	same as above except replace(NP1,random NP) <b>or</b> , replace(VB,random VB)	<i>People are playing baseball.</i> <i>People are playing baseball.</i>	<i>Are cats playing baseball?</i> <i>Are people making coffee?</i>
number	EX+is/are+num+NP→How many+NP+is/are+EX? num+NP+VB→replace(num,how many)	<i>There are two cats.</i> <i>Six cars are parked.</i>	<i>How many cats are there?</i> <i>How many cars are parked?</i>
others	NP1(+MD)+VB → WP(+MD)+VB? <b>or</b> , →WP+MD/conjug(“do/be”,tense(VB))+NP+“do/doing”? obj/property+VB(+NP) → WP/WDT+category+VP(+NP)	<i>A boy is running.</i> <i>A boy is running.</i> <i>An apple is shown.</i>	<i>Who is running?</i> <i>What is a boy doing?</i> <i>What fruit is shown?</i>
	NP1+VB(+NP2)+IN+NP3 →WRB+conjug(“do/be”,tense(VB))+NP1+conjug(VB,tense(VB))?	– <i>Dogs run in a park.</i>	– <i>Where do dogs run?</i>

## Converting Captions

We also provide an augmented version of the dataset by converting the human-written captions for images into questions and answers. Apart from yes/no questions, the answers to the generated questions are the captions themselves, eliminating the burden for generating reliable answers. Most images in MS COCO come with 5 captions, and we generated distinct question for each caption, whose conversion rule is shown in Table 2.

We assigned at least two yes/no questions with one “yes” and one “no” to all images, roughly balancing the number of answers with “yes,” and “no.” Questions with affirmative answers involving “yes” were generated by simply rephrasing the caption such that the question asks to confirm the contents in the caption, for which the answer is an affirmative statement of the question (which is the caption itself),

accompanied by “Yes,” in the beginning.

Questions with non-affirmative answers involving “no” were generated by substituting parts of captions with random actions or agents. For agents, we randomly choose one class from 1,000 object classes of ILSVRC 2014 object detection task, and substitute it for given agent in the caption. For actions, we randomly choose one class from 101 action classes of UCF-101 (Soomro, Zamir, and Shah 2012) plus 20 classes of activities of daily living (ADL) from (Ohnishi et al. 2016), and substitute it for given verb phrase. Resulting questions are frequently of interesting non-sensical type, such as “are the birds doing push-ups on the tree?,” for which the answer is simply a negation of the question, accompanied by “No,” in the beginning. We expect that such set of questions and answers can also potentially help the learning of distinction between common sense and nonsense.

Table 3: Statistics for VQA and two versions of FSVQA dataset. Note that statistics for VQA dataset were computed from the most frequent answers for each question.

Dataset	split	number of Q&A pairs	avg len answer	number of unique answers	number of unique words	top 1k answers coverage(%)
VQA (Antol et al. 2015)	train	248,349	1.11	17,089	11,673	86.74
	val	121,512	1.12	10,922	8,078	86.56
	all	369,861	1.11	22,497	14,528	86.56
FSVQA	train	248,349	6.06	173,016	32,790	13.02
	val	121,512	6.04	92,135	22,199	13.27
	all	369,861	6.05	248,223	39,663	12.66
FSVQA aug.	train	662,462	10.38	579,761	57,468	4.98
	val	324,166	10.22	292,525	39,344	5.05
	all	986,628	10.33	853,001	69,486	4.83

Table 4: Number of unique answers for each category.

Dataset	yes/no	number	others
VQA	2	817	21,678
FSVQA	98,700	27,119	122,324
FSVQA aug.	485,428	72,970	294,603

We also generated questions that are category-specific. When an object or a property of a specific category is referred to, we replace it with the category name, preceded by *wh*-determiner or *wh*-pronoun, to ask which object or property it belongs to. Our manually pre-defined categories included color, animal, room, food, transportation, and sport. Other rules are mostly reversed process of the conversion rules in Table 1, in which we deliberately mask certain content of the caption, replace it with appropriate question forms, and reorder the sentence.

## Statistics

Table 3 shows statistics for original VQA dataset and two versions of FSVQA dataset. Both versions of FSVQA contain much longer answers on average, and the number of unique answers is more than ten times larger in the regular version and about 38 times larger in the augmented version compared to VQA dataset. Augmented version is longer since captions in MS COCO tend to be longer than questions in VQA. The size of vocabularies is also many times larger in both versions. Note that, consistently with the conversion process, only the most frequent answer from 10 answers per question was taken into consideration for VQA dataset’s statistics.

Comparing the coverage of 1,000 most frequent answers in the dataset shows even more striking contrast. In VQA dataset, 1,000 most frequent answers covered about 86.5% of the entire dataset, so that learning a small set of frequent answers could perform reasonably well. In FSVQA, the coverage by 1,000 most frequent answers is merely 12.7% and 4.8% for regular and augmented version respectively, which is clearly much less than VQA dataset. Thus, it adds a critical amount of computational complexity, in which less frequent answers cannot easily be disregarded.

Table 4 shows the number of unique answers for each category. While FSVQA has much more answers in all cate-

gories, most striking example is in the yes/no category. VQA dataset essentially contains only two answers “yes” and “no” for yes/no questions (infrequent answers such as “not sure” were filtered out in the conversion process). In fact, answering with “yes” alone for all questions achieved 70.97% accuracy for yes/no category in the original VQA dataset. On the contrary, FSVQA datasets contain approximately 49,000 times more answer and 240,000 times more answers for yes/no category in each version. It becomes clear again that FSVQA cannot be taken advantage of by manipulating a small set of frequent answers.

Figure 2 shows the percentage distribution of answers in the respective datasets for number of words in the answers. While over 90% of the answers in VQA dataset are single words, both versions of FSVQA dataset show much smoother distribution over a wide range of number of words.

## Experiment

### Setting

We used 4096-dimensional features from the second fully-connected layer (fc7) of VGG (Simonyan and Zisserman 2014) with 19 layers, trained on ImageNet (Deng et al. 2009), as our image features. Words in the question were input to LSTM (Hochreiter and Schmidhuber 1997) one at a time as one-hot vector, where the dictionary contains only the words appearing more than once. Image features and question features are then mapped to common embedding space as a 1,024-dimensional vector. Batch size was 500 and training was performed for 300 epochs.

We trained only with the answers that appear twice or more in train split, as using all unique answers in the dataset fails to run, with required memory far beyond the capacity of most of the contemporary GPUs, NVIDIA Tesla 40m in our case. 20,130 answers appear more than once in regular version, covering 95,340 questions from 62,292 images, and 23,400 answers appear more than once in the augmented version, which cover 105,563 questions from 64,060 images. This is only about 25% and 15.5% of the entire train split in respective version, which again shows a striking contrast with the original VQA dataset, in which only 1,000 answers covered up to 86.5% of the dataset.

Following (Antol et al. 2015), we examined the effect of

Table 5: Performances of the generated answers on evaluation metrics

Dataset	Model	VQA Acc.	FSVQA Acc.	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	Cider
FSVQA	LSTM Q+I	32.54	16.74	56.9	41.5	31.8	23.9	0.233	2.904
	LSTM Q	31.00	16.18	54.2	39.3	30.4	23.3	0.219	2.755
	Image	9.72	1.35	15.4	2.6	0.8	0.2	0.101	0.412
FSVQA aug.	LSTM Q+I	32.57	7.09	42.2	23.7	15.1	10.2	0.150	1.798
	LSTM Q	31.53	6.69	40.5	22.9	14.6	10.0	0.141	1.780
	Image	14.41	0.87	26.3	8.7	4.6	3.2	0.094	0.601




Figure 2: Distribution of answer lengths as the number of words over each dataset.

removing images, since questions alone may frequently provide sufficient amount of clue to correct answers. Training procedure is identical as above, except only question features are mapped to the common embedding space since there are no image features.

Conversely, we also examined an approach where only image features are concerned. This requires a slightly different training procedure, as it does not involve a series of one-hot vector inputs. We followed the conventional approach used in image captioning task (Vinyals et al. 2015), where the image features are fixed, and a stack of LSTM units learns to generate the ground truth captions. Each LSTM unit generates one word at a time, which in turn enters the next LSTM unit. The only difference in our case is that the ground truth captions are replaced by full-sentence answers.

## Evaluation

**Metrics** Evaluating the results from the experiments also poses a challenge. Since original VQA dataset consisted mostly of short answers, evaluation was as simple as matching the results with ground truths, and yielding the percentage. Yet, since we now have full-sentence answers, simply matching the results with ground truths will not be compatible. For example, “yes, the color of the cat is red” for the question in which the original ground truth is “yes,” will be classified as incorrect using the current evaluation tool for original VQA.

We thus come up with a set of mutually complementary ways of evaluating full-sentence answers. First, we employ the frequently used evaluation metrics for image captioning task, namely BLEU (Papineni et al. 2002), METEOR (Denkowski and Lavie 2014), and CIDEr (Vedantam, Zitnick, and Parikh 2015). The goal is to quantify the overall resemblance of the results to the ground truth answers.

However, higher performance on these evaluation metrics does not necessarily imply that it is a more accurate answer. For instance, “the color of the car is red” will have higher scores than “it is blue” for a ground-truth answer “the color of the car is blue,” due to more tokens being identical. Yet, the former is clearly an incorrect answer, whereas the latter should be considered correct. In order to overcome this drawback, we also employ a simple complementary evaluation metric, whose procedure is as follows: we examine whether the short answer from the original dataset is present in the generated result, and extract the short answer if present. If not, we leave the answer blank. Extracted terms in this way are tested with the evaluation tool for original VQA. Using the previous example, the rationale is that as long as the original short answer “blue” is present in the generated result, it can be assumed that the answer is correct. We refer to this metric as VQA accuracy. Note that, for augmented version, generated answers for only the original subset are extracted to measure VQA accuracy, since there are no ground truth VQA answers for the augmented segment.

However, there also exist cases in which VQA accuracy can be misleading, since the rest of the context may not be compatible with the question. For example, “yes, the color is blue.” will be considered correct if the original answer is “yes,” but it should not be considered correct if the question was “is it raining?” In fact, this was one of the underlying concerns in the original VQA dataset, since we cannot be sure whether “yes” or “no” was generated in the right sense or purely by chance. In order to alleviate this issue, we also report FSVQA accuracy, which is the percentage in which the ground truth answer in FSVQA dataset contains the generated answer. Since the answers have to be matched at the sentence level, it assures us with high confidence that the answer was correct in the intended context.

**Results & Discussion** Results for all metrics are shown in Table 5. Note that evaluation is performed on the results for validation split, since ground truths for test split are not publicly available. While each metric is concerned with slightly different aspect of the answers, results shows that they gen-





Q: How many cows are there?  
VQA: 5  
Q+I: There are 5 cows.  
Q-only: There are 3 cows.  
I-only: This is cow.



Q: What sport is being played?  
VQA: tennis  
Q+I: Tennis is being played.  
Q-only: Baseball is being played.  
I-only: This is tennis.



Q: What kind of balls is the boy holding?  
VQA: tennis  
Q+I: This person is holding tennis racket.  
Q-only: The boy is holding wii.  
I-only: This is tennis.



Q: Is this in a park?  
VQA: yes  
Q+I: Yes, this is in a park.  
Q-only: Yes, this is in a park.  
I-only: Yes, it is daytime.



Q: Is there a rug in the bedroom?  
VQA: yes  
Q+I: Yes, there is a rug on the floor.  
Q-only: Yes, there is a chair in the picture.  
I-only: This is kitchen.



Q: Are they sleeping?  
VQA: no  
Q+I: No, they are not sleeping.  
Q-only: No, they are not sleeping.  
I-only: Yes, the bed is made.

Figure 3: Examples of questions and generated answers by each baseline approach, along with ground truth answer from the original VQA dataset.

erally tend to agree with each other. Figure 3 shows examples of generated full-sentence answers for each model, along with the ground truth answer from the original VQA dataset.

As expected, answers generated from using both question and image features turn out to be most reliable. Answers from question features alone result in answers that match the questions but are frequently out of visual context given by the image. Likewise, answers generated from image features alone fit the images but are frequently out of textual context given by the question. It is notable that using image features alone performs very poorly, whereas using question features alone results in performances comparable to using both features. One plausible explanation is that, since using image features alone always generates the same answer for the same image regardless of the question, it can only get 1 out of  $k$  questions correctly at best, where  $k$  is the number of questions per image. On the contrary, using question features alone essentially reduces the problem to a semantic Q&A task, which can be handled one at a time. This tendency is consistent with the results reported in (Antol et al. 2015). It must nevertheless be reminded that the best performances in both (Antol et al. 2015) and our experiment were achieved with the presence of both visual and textual clues.

## Conclusion

We introduced FSVQA, a publicly available dataset consisting of nearly 1 million pairs of questions and full-sentence answers for images, built by applying linguistic rules to existing datasets. While pushing forward the VQA task to a more human-like stage, it poses many extra complexities. We examined baseline approaches for tackling this novel task. Applying some of the successful approaches from the original VQA task, such as attention mechanism, will be an intriguing and important future work. Whether generative approach can play more role in the future, as the number of answers grows larger and classification approach becomes less efficient, is also of interest. We invite the research community to come up with an innovative and efficient way to improve the performance on FSVQA.

## Acknowledgement

This work was funded by ImPACT Program of Council for Science, Technology and Innovation (Cabinet Office, Government of Japan).

## References

- [Antol et al. 2015] Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Lawrence Zitnick, C.; and Parikh, D. 2015. Vqa: Visual question answering. In *ICCV*.
- [Deng et al. 2009] Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*.
- [Denkowski and Lavie 2014] Denkowski, M., and Lavie, A. 2014. Meteor Universal: Language Specific Translation Evaluation for Any Target Language. In *EACL Workshop on Statistical Machine Translation*.
- [Fukui et al. 2016] Fukui, A.; Park, D.; Yang, D.; Rohrbach, A.; Darrell, T.; and Rohrbach, M. 2016. Multimodal Compact Bilinear Pooling for Visual Question Answering and Visual Grounding. In *EMNLP*.
- [Hochreiter and Schmidhuber 1997] Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural Computation* 9(8):1735–1780.
- [Karpathy and Li 2015] Karpathy, A., and Li, F. 2015. Deep Visual-Semantic Alignments for Generating Image Descriptions. In *CVPR*.
- [Kim et al. 2016] Kim, J.; Lee, S.; Kwak, D.; Heo, M.; Kim, J.; Ha, J.; and Zhang, B. 2016. Multimodal Residual Learning for Visual QA. In *NIPS*.
- [Lin et al. 2014] Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollr, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common Objects in Context. In *ECCV*.
- [Lu et al. 2016] Lu, J.; Yang, J.; Batra, D.; and Parikh, D. 2016. Hierarchical Question-Image Co-Attention for Visual Question Answering. In *NIPS*.
- [Malinowski and Fritz 2014] Malinowski, M., and Fritz, M. 2014. Towards A Visual Turing Challenge. In *NIPS Workshop on Learning Semantics*.
- [Marcus et al. 1994] Marcus, M.; Kim, G.; Marcinkiewicz, M.; MacIntyre, R.; Bies, A.; Ferguson, M.; Katz, K.; and Schasberger, B. 1994. The Penn Treebank: annotating predicate argument structure. In *HLT*.
- [Mei et al. 2008] Mei, T.; Wang, Y.; Hua, X.; Gong, S.; and Li, S. 2008. Coherent Image Annotation by Learning Semantic Distance. In *CVPR*.
- [Noh and Han 2016] Noh, H., and Han, B. 2016. Training recurrent answering units with joint loss minimization for vqa. *arXiv:1606.3647*.
- [Ohnishi et al. 2016] Ohnishi, K.; Kanehira, A.; Kanazaki, A.; and Harada, T. 2016. Recognizing activities of daily living with a wrist-mounted camera. In *CVPR*.
- [Papineni et al. 2002] Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W. 2002. BLEU: A Method for Automatic Evaluation of Machine Translation. In *ACL*.
- [Ren, Kiros, and Zemel 2015] Ren, M.; Kiros, R.; and Zemel, R. 2015. Exploring Models and Data for Image Question Answering. In *NIPS*.
- [Ritter, Cherry, and Dolan 2011] Ritter, A.; Cherry, C.; and Dolan, B. 2011. Data-Driven Response Generation in Social Media. In *EMNLP*.
- [Saito et al. 2016] Saito, K.; Shin, A.; Ushiku, Y.; and Harada, T. 2016. Dualnet: Domain-invariant network for visual question answering. *arXiv:1606.6108*.
- [Simonyan and Zisserman 2014] Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*.
- [Soomro, Zamir, and Shah 2012] Soomro, K.; Zamir, A.; and Shah, M. 2012. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv:1212.0402*.
- [Vedantam, Zitnick, and Parikh 2015] Vedantam, R.; Zitnick, C. L.; and Parikh, D. 2015. CIDEr: Consensus-based Image Description Evaluation. In *CVPR*.
- [Vinyals et al. 2015] Vinyals, O.; Toshev, A.; Bengio, S.; and Erhan, D. 2015. Show and tell: A neural image caption generator. In *CVPR*.
- [Wang et al. 2009] Wang, C.; Yan, S.; Zhang, L.; and Zhang, H. 2009. Multi-Label Sparse Coding for Automatic Image Annotation. In *CVPR*.
- [Xu et al. 2015] Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhudinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In *ICML*.
- [Yang et al. 2016] Yang, Z.; He, X.; Gao, J.; Deng, L.; and Smola, A. 2016. Stacked Attention Networks for Image Question Answering. In *CVPR*.